NASA/TM—2000-209803

An Assessment of Gigabit Ethernet Technology and Its Applications at the NASA Glenn Research Center

Catherine Murphy Bakes
Kent Sate University, Kent, Ohio

Chan M. Kim and Calvin T. Ramos
Glenn Research Center, Cleveland, Ohio

February 2000

The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the Lead Center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized data bases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA STI Program Home Page at ***http://www.sti.nasa.gov***

- E-mail your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Telephone the NASA Access Help Desk at (301) 621-0390

- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076

NASA/TM—2000-209803

An Assessment of Gigabit Ethernet Technology and Its Applications at the NASA Glenn Research Center

Catherine Murphy Bakes
Kent Sate University, Kent, Ohio

Chan M. Kim and Calvin T. Ramos
Glenn Research Center, Cleveland, Ohio

National Aeronautics and
Space Administration

Glenn Research Center

February 2000

Acknowledgments

The authors would like to thank David Pleva and Raymond Schutz for their active participation in many discussions regarding Gigabit Ethernet and for their helpful comments and suggestions. This research was supported by the High Performance Computing and Communication Project (HPCCP) at the NASA Glenn Research Center.

Available from

NASA Center for Aerospace Information
7121 Standard Drive
Hanover, MD 21076
Price Code: A03

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22100
Price Code: A03

An Assessment of Gigabit Ethernet Technology and Its Applications at the NASA Glenn Research Center

Catherine Murphy Bakes
Kent State University
Department of Administrative Sciences
College of Business Administration
Kent, Ohio 44242
Phone: (330) 672–2750 Ext. 338, Fax: (330) 672–2448
Email: cbakes@bsa3.kent.edu

Chan M. Kim and Calvin T. Ramos
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

Abstract: This paper describes Gigabit Ethernet and its role in supporting R&D programs at NASA Glenn. These programs require an advanced high-speed network capable of transporting multimedia traffic, including real-time visualization, high-resolution graphics, and scientific data. GigE is a 1 Gbps extension to 10 and 100 Mbps Ethernet. The IEEE 802.3z and 802.3ab standards define the MAC layer and 1000BASE-X and 1000BASE-T physical layer specifications for GigE. GigE switches and buffered distributors support IEEE 802.3x flow control. The paper also compares GigE with ATM in terms of quality of service, data rate, throughput, scalability, interoperability, network management, and cost of ownership.

1. Introduction

As application integration, distribution, and collaboration have become more common in the network computing environment, the demand for better, faster networking technologies has grown. Not only do many of today's emerging applications demand higher bandwidth, but they also require better service in terms of quality, reliability, and security. Only a few years ago, many campus network infrastructures consisted of shared 10 Mbps Ethernet technologies and employed software-based routing. The majority of the traffic stayed within a department's local boundary and typical applications, such as electronic mail, file transfers, and printing, cooperatively shared the available bandwidth. With the development of Application Specific Integrated Circuit (ASIC) technology and communication protocol standards, high performance switches with improved services have become the most popular solution to meet new challenges in network computing. 10/100 Mbps switched Ethernet/Fast Ethernet is replacing concentrator-based 10Base-T shared Ethernet and Layer 3 switches boast wire-speed routing capabilities. Also, Asynchronous Transfer Mode (ATM) has brought a new paradigm in networking technology through its OC-3/OC-12 speeds, voice-data integration, and superb quality of service. On the Ethernet side, 1000 Mbps Gigabit Ethernet technology has made substantial progress in many key standards issues and, due to its relative simplicity and low-cost, has gained favorable support from network professionals as a formidable choice for the next generation of Local Area Network (LAN) technology.

This paper deals with an overview of Gigabit Ethernet network technology and its role in supporting research and development activities at the NASA Glenn Research Center. NASA's current mission revolves around four strategic enterprises, i.e., space science, mission to planet earth, human exploration and development of space, and aeronautics and space transportation technology. The NASA Glenn Research Center is actively involved in many of these enterprises through in-house research and collaboration with various governmental agencies, universities, and industry (Bakes, et al., 1996).

In order to support NASA Glenn's scientific missions, both computing and networking infrastructures have to provide adequate speed, quality, and reliability. The computing resources at NASA Glenn include desktop workstations and servers based on either Windows or Unix operating systems, plus a few LINUX systems that have been deployed recently. For computationally intensive scientific and engineering applications, a cluster of servers with multiple processors is utilized. A typical desktop workstation, in order to run necessary applications, is connected to servers through a high speed network. In addition to scientific and engineering applications, many mission-critical applications, such as file transfers, back-ups, web access, and electronic mail, occupy a large portion of the network's traffic. At the heart of all these applications is the campus backbone network, which provides a variety of services to reliably deliver high throughput transport of critical data.

Optical fiber is the transmission medium used for the campus backbone that carries local traffic between buildings at NASA Glenn. This network includes many strands of single mode and multimode fiber and allows any attached station to access a potential bandwidth ranging from 100s of Mbps to multiple Gbps. The backbone interfaces to various fiber, coaxial cable, and copper local area networks within buildings and has been extended to selected servers and researchers' workstations that require high-speed connectivity. To further exploit the benefits of fiber, the public Wide Area Network (WAN) providers have installed fiber throughout the entire geographic area of interest to the Glenn community. This enables Glenn researchers to access remote sites from their desktop workstations and enables users at remote sites to access Glenn facilities via a seamless fiber network at data rates up to and beyond 155 Mbps. Future network initiatives will include Gbps speeds, especially for backbones.

The NASA Glenn Research Center has active programs in computational aerodynamics, material science, structure dynamics, space communication, and space sciences. These programs require a communications network capable of transporting multimedia traffic, real time visualization, and data collected from scientific experiments. The diverse tasks performed by powerful desktop workstations, local clusters, and central servers also place a large demand on the network. In addition, NASA Glenn's research community is developing the next generation of computing applications and exploring their network implications. Intelligent Synthesis Environment (ISE) is an ambitious program to develop and implement tools and processes that enable geographically dispersed scientists, technologists, and engineers, with diverse expertise and interests, to function as a coherent team in the conceptualization, design, development, and execution of NASA's missions. The new paradigm of Agency computing initiatives is to be supported by technologies such as multimedia desktop conferencing, distributed object technology, and web-database integration.

In order to support and perform Glenn's mission-critical applications, desktop workstations are typically equipped with 10Base-T Ethernet connections to servers and the backbone. Although Ethernet is still the most popular LAN technology in use today at the desktop, the bandwidth offered by Ethernet becomes inadequate for acceptable performance as the sheer volume of network traffic increases. Fast Ethernet, or 100Base-T Ethernet, technology has provided a smooth evolution from 10 Mbps to 100 Mbps performance and has been adopted for server-to-server communications. The demand for higher bandwidth to the desktop has also grown for many end users and led to a need for an even higher speed network technology at the backbone. Gigabit Ethernet technology provides 1000 Mbps bandwidth at lower cost than other technologies of comparable speed and is thus a natural upgrade strategy for legacy Ethernets. It is an extension and enhancement to Ethernet and Fast Ethernet that offers scalability and 10 times the performance of Fast Ethernet at two to three times the cost. It employs the same Carrier Sense, Multiple Access/Collision Detection (CSMA/CD) protocol, frame format, and frame size as its predecessors. As a result, many existing networks can be extended to gigabit speeds, at reasonable initial cost and without re-educating support staff or investing in additional protocol stacks or middleware.

The deployment of Gigabit Ethernet technology at NASA Glenn will enable the development of many bandwidth-intensive, interdisciplinary applications. The combination of high speed, standards-based Quality of Service (QoS) features, fast routing, and ease of management makes Gigabit Ethernet an ideal solution for next generation network technology.

2. Gigabit Ethernet Standards

In July 1996, the IEEE 802.3 working group created the IEEE 802.3z Gigabit Ethernet (GigE) task force with the objective of developing a GigE standard. The IEEE 802.3z standard, a 1 Gbps, backward compatible, extension to the IEEE 802.3 standards for 10 and 100 Mbps Ethernet, was completed in 1998. Like 10 and 100 Mbps Ethernet, it is a data link and physical layer technology only.

2.1 Gigabit Ethernet Media Access Control Layer

Gigabit Ethernet uses the same frame format as its 10 and 100 Mbps predecessors, with frames of 64 to 1,518 bytes, excluding preamble and Start-of-Frame Delimiter (SFD), and a 96 bit Inter-Frame Gap (IFG). Figure 1 shows the basic IEEE 802.3 frame format.

7	1	6	6	2	46-1500	4
Preamble	SFD	Destination address	Source address	Type/ Length	Data + Pad	CRC

<------------------------ 64 byte minimum ------------------------>

Figure 1.—IEEE 802.3 frame format.

All three Ethernet speeds are able to operate in Half-Duplex (HDX) mode for shared-media LANs and in Full-Duplex (FDX) mode for dedicated, switched connections. Shared Ethernet networks use the IEEE 802.3 CSMA/CD access method to resolve contention. The CSMA/CD algorithm requires a sender to listen to the network before transmitting a frame to determine if the channel is free (i.e., carrier sense), and to continue listening while transmitting to determine if the frame experiences a collision (IEEE 802.3, 1998). A sender that detects a collision performs a jam, backoff, and reschedule sequence during which it stops sending and generates a backoff interval that schedules the next transmission attempt. The backoff interval is the product of the Ethernet slot time, which equals the round trip propagation delay on a network of maximum size, and a randomly generated integer. On a correctly configured CSMA/CD network, a sending station must be able to detect a collision before it completes transmission of a colliding frame and the maximum time, from the start of the transmission until the sender detects the collision, must be less than the Ethernet slot time (Kadambi et al., 1998). This requires the time to transmit a minimum length frame to be longer than the round trip propagation delay of the network. On an oversized network (i.e., one where the round-trip delay is longer than the slot time), it is possible for a station to complete transmission before detecting a collision. This late collision problem violates the CSMA/CD access method and could lead to network instability.

When the IEEE 802.3u 100BASE-T standard, known as Fast Ethernet, was approved in 1995, it represented a tenfold increase in data rate to 100 Mbps, from the 10BASE-T rate of 10 Mbps, and a corresponding tenfold reduction in the time to transmit a minimum length (i.e., 64 byte) frame. In order to prevent late collisions when two stations simultaneously transmit minimum length frames, the IEEE 802.3u task force considered decreasing the maximum network diameter, or increasing the minimum frame length. The task force decided to decrease the maximum diameter for 100BASE-T LANs to the order of 200 m with Category 5 unshielded twisted-pair (412 meters with multimode fiber) and to leave the minimum frame length unchanged at 64 bytes (IEEE 802.3, 1998; Seifert, 1998).

Gigabit Ethernet represents another tenfold increase in data rate and, again, reduces the time required to transmit a frame by a factor of 10. This could have been achieved by a further reduction in network diameter. However, since a 20-meter network diameter was considered impractical, the IEEE 802.3z working committee, in essence, redefined the MAC layer for GigE by adding a mechanism to make a 200-meter network diameter possible at 1 Gbps. This mechanism is known as “carrier extension” (IEEE 802.3, 1998)**.** Whenever a shared gigabit network adapter transmits a frame shorter than 512 bytes long, it adds a new carrier extension field of up to 448 bytes (3,584 bits) to the frame and continues to monitor for collisions while sending this special signal. The carrier extension field follows the CRC field and contains a sequence of special non-data “extended carrier” symbols that are not considered part of the frame. The CRC remainder is calculated only on the original frame (i.e., without extension symbols) and the frame plus carrier extension lasts for a minimum of 512 bytes. To prevent late collisions for the required distance limits, GigE also extends the Ethernet slot time to 512 bytes (4,096 bits), from 64 bytes (512 bits) for Ethernet and Fast Ethernet. The 64-byte minimum frame length and 96 bit IFG have not changed for GigE and frames longer than 512 bytes are not extended. Figure 2 shows the Gigabit Ethernet frame format when Carrier Extension is used.

8	6	6	2	46-1500	4	0-448
Preamble + SFD	Destination address	Source address	Type/ Length	Data + Pad	CRC	Carrier Extension

<----------------------------- 512 byte minimum ------------------------------->

Figure 2.—Format of Gigabit Ethernet frame with carrier extension.

While Gigabit Ethernet should be able to transmit 640 bytes (rather than 512 bytes) in the time it takes a Fast Ethernet interface to transmit 64 bytes, the IEEE 802.3z working committee decided that a 640-byte extension was too inefficient and shortened the extension to 512 bytes. To facilitate the shorter extension, they reduced the number of repeater hops to one, from two permitted in 100BASE-T, and basically eliminated the safety margin built into engineering specifications for earlier Ethernet implementations. Consequently, to avoid instability on GigE networks, manufacturers must strictly adhere to timing specifications.

In order to utilize available bandwidth more efficiently, and to offset the adverse impact of carrier extension and slot time extension on performance when traffic consists of short frames, an optional new feature has been incorporated into the CSMA/CD algorithm for shared GigE networks. This feature, called “frame bursting,” allows a gigabit network adapter to transmit multiple short frames with a single arbitration for the channel (Cunningham et al., 1999; IEEE 802.3, 1998; Seifert, 1998). When a station that has been idle transmits a frame, which may or may not require carrier extension, a “burst timer” is started. If this first frame is successfully transmitted (i.e., without a collision), then the sending station has the option of transmitting additional frames, subject to the conditions that it has another frame in its transmit queue and the burst timer has not expired. Since the sender will have acquired the medium by the time it completes transmission of the first frame, subsequent frames within a burst are guaranteed not to experience collisions and do not need to be extended. The start of the last frame in a burst must occur before the burst timer expires, but its transmission may extend beyond the burst timer limit of 8,192 bytes (65,536 bits). Thus, the maximum duration of a single transmission can be the sum of the burst length and maximum frame size, which is more than 6 times the maximum frame size. A 96 bit interframe gap is transmitted after each frame in the burst. Depending on the degree of traffic burstiness, waiting for a sending station to complete transmission of a burst of frames could increase the delays experienced by other stations that have frames ready for transmission.

On switched GigE networks, switches transmit and receive data on different fibers (or wire pairs) in point-to-point configurations and never experience collisions. Therefore, they do not use CSMA/CD, carrier extension (they use the regular 64 byte minimum frame size), slot time extension, or frame bursting, which are used with shared HDX GigE. Also, eliminating collisions removes the CSMA/CD timing restriction that limits the maximum diameter for shared networks (Kadambi et al., 1998). Longer distances thus become possible for switched links because they are based on link characteristics such as attenuation rather than on propagation delay. Other than increasing the bit rate to 1 Gbps, no MAC layer changes from switched Fast Ethernet were required for switched GigE, which almost always runs in FDX mode.

In order to improve performance on server farms and computer clusters, some Gigabit Ethernet vendors implement proprietary “jumbo frames,” which are typically between 9 and 64 kilobytes in length. Jumbo frames require less processing than shorter frames, thereby freeing server CPUs for other tasks. However, they do not conform to Ethernet’s 1,518-byte limit and their use and availability are limited.

2.2 Gigabit Ethernet Physical Layer Standards

Gigabit Ethernet supports four different physical layer implementations, three of which are defined in the IEEE 802.3z standard (IEEE 802.3, 1998). The fourth is defined in the IEEE 802.3ab standard, to be discussed later in this section (IEEE 802.3ab, 1999). IEEE 802.3z provides the specifications for the 1000BASE-LX, 1000BASE-SX, and 1000BASE-CX physical layers which, together, are generically referred to as 1000BASE-X. 1000BASE-LX networks support three types of optical fiber and 1000BASE-SX networks support 2 types. 1000BASE-LX can operate over a pair of 10 μm (core diameter) Single Mode Fibers (SMF), or 50 or 62.5 μm Multimode Fibers (MMF), and uses long-wavelength (1300 nm nominal) lasers. 1000BASE-SX specifies operation over a pair of 50 or 62.5 μm multimode optical fibers and uses short-wavelength (850 nm nominal) lasers. Since fiber is typically more expensive than copper, especially in terms of termination and installation, IEEE 802.3z also includes the 1000BASE-CX specification that operates over two pairs of 150 ohm shielded, balanced, copper cable. For switched links, 1000BASE-LX supports distances up to 5 km with SMF and up to 550 meters with both types of MMF. Depending on the modal bandwidth, 1000BASE-SX support maximum distances ranging from 500 to 550 meters over 50 μm MMF and from 220 to 275 meters over 62.5 μm MMF. Due to the timing constraints imposed by CSMA/CD, the maximum distance for a shared GigE fiber link is limited to 110 meters. Maximum distances on 1000BASE-CX networks are limited to 25 meters for both shared and switched links. This is due to the characteristics of the physical medium itself, and not to the constraints imposed by a CSMA/CD collision domain.

The IEEE 802.3z task force extensively tested the operating characteristics of lasers on multimode fiber. These tests revealed a jitter component caused by a condition known as Differential Mode Delay (DMD) that occurs in certain MMF fibers when using laser diodes (Gigabit Ethernet Alliance, May 1999; IEEE 802.3, 1998; Seifert, 1998). DMD is a phenomenon in which light rays in a MMF travel through several paths of different lengths, so that pulses launched at one end of a fiber spread out in space and time, causing poor signal reception at the receiving end. The solution developed by the IEEE 802.3z task force is called a Conditioned Launch (CL). A CL spreads out the laser light-source output so that it looks like an LED source for which the cable was designed. By spreading the power across the core, more or less equally in all modes, the DMD effect can be minimized. DMD does not arise in SMF because there is only one ray or propagation mode of the light signal.

The IEEE 802.3z task force drew heavily from the physical layer developed by ANSI for the X3.230-1994 Fibre Channel standard, which is a technology for interconnecting workstations, supercomputers, storage devices and peripherals at gigabit speeds. Fibre Channel uses an 8B/10B coding scheme which encodes each 8 bits of data into a 10 bit "code group." 1000BASE-X networks are also based on 8B/10B coding and they use a signaling rate of 1.25 Gbaud to achieve the 1 Gbps date rate.

Gigabit Ethernet divides the OSI physical layer into 4 sublayers and 2 interfaces (Figure 3) (IEEE 802.3, 1998; IEEE 802.3ab, 1999; Kadambi et al., 1998). The Reconciliation Sublayer (RS) and optional Gigabit Media Independent Interface (GMII) are common to all GigE media types. The remaining three sublayers, the Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and Physical Medium Dependent (PMD) sublayer, and the Medium Dependent Interface (MDI) are dependent on the particular physical media and data encoding method. RS maps the bit serial MAC interface to the multiple bit wide data path defined by GMII. GMII, which is the 1 Gbps equivalent to the 100 Mbps Media Independent Interface (MII), provides a logical signal interface between the GigE MAC and physical layers and allows the MAC layer to be connected to different cable types. It is generally not used with 1000BASE-X since all 1000BASE-X media types use the same 8B/10B-encoding scheme. PCS provides data coding and decoding functions and, for shared operation, it also generates Carrier Sense and Collision Detect indications. PMA defines a mechanism for converting code groups to and from a serial stream, which it passes to PCS. The 1000BASE-X PCS uses 8B/10B encoding and the 1000BASE-X PMA sublayer serializes 10-bit code groups before transmission and deserializes a received stream into code groups. These two sublayers are common to all three 1000BASE-X PMDs. PMD defines the physical layer signaling used for various media and converts a serial bit stream from PMA into a signal appropriate for the specific physical media.

Each GigE media type requires a corresponding PMD. The LX and SX PMDs provide the specifications for the various optical fiber media (i.e., SMF and 50 and 62.5 μm MMF) and optical wavelengths (i.e., 1270-1355 nm and 770-860 nm) supported by GigE. They perform electrical to optical conversions for serial bit streams from the PMA sublayer, and vice versa. The CX PMD provides specifications for 2 pair, shielded copper cable, along with the necessary line drivers, receivers, and system signal budgets. The MDI, which is a part of PMD, defines the connectors for the different media types.

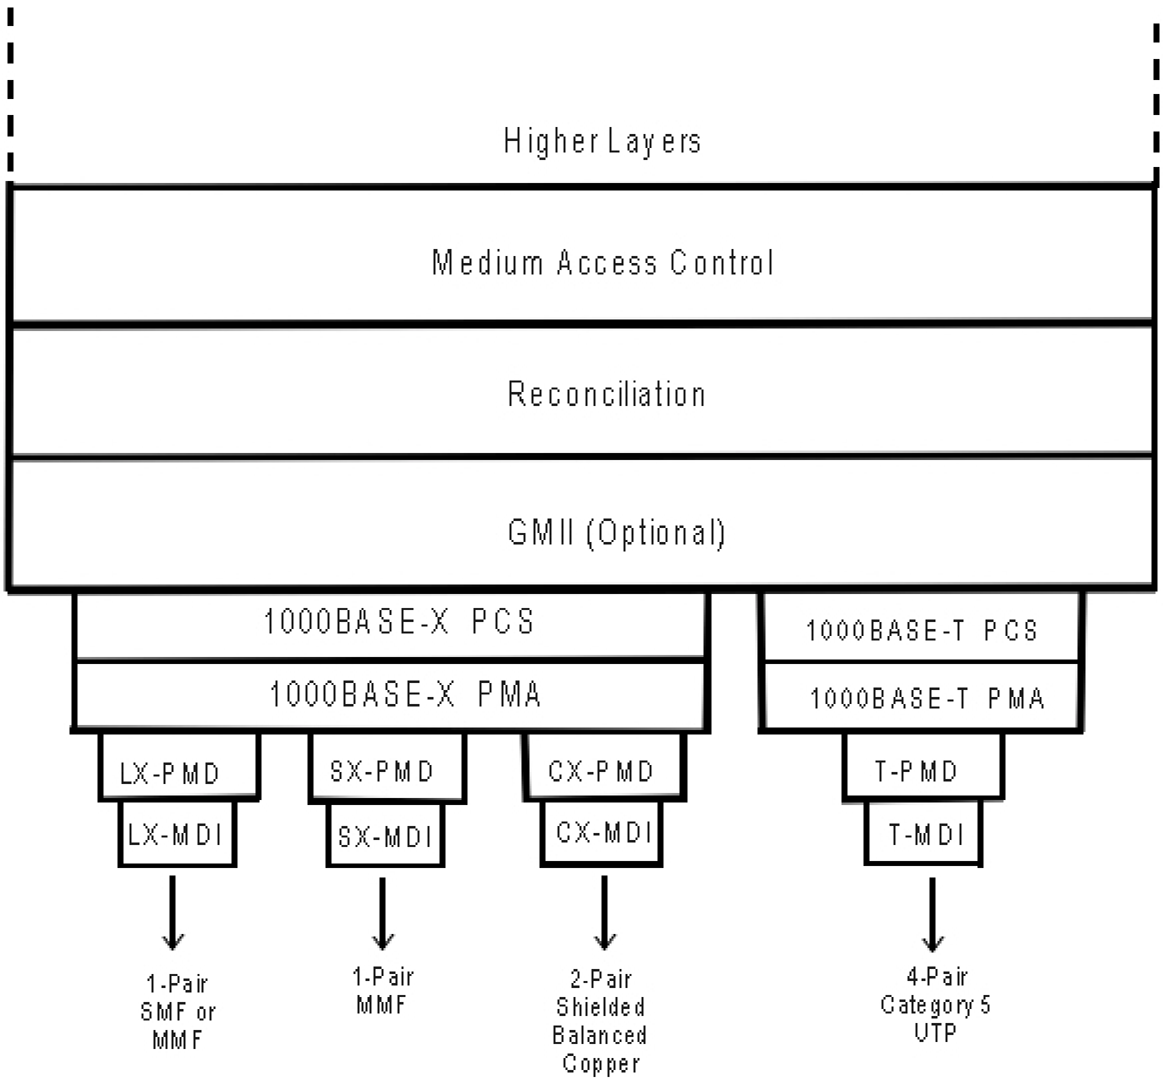

Figure 3.—Gigabit Ethernet reference model.

A separate but related Gigabit Ethernet task force, IEEE 802.3ab, developed the 1000BASE-T standard, which was completed in June 1999 (Gigabit Ethernet Alliance, 1997; IEEE 802.3ab, 1999). 1000BASE-T adds a fourth physical layer, with a different set of PCS/PMA/PMD sublayer specifications, to the basic technology defined by 1000BASE-X. It defines 1 Gbps operation over four pairs of category 5 (or better) UTP cable, supports shared and switched link distances of up to 100 meters, and is capable of sending and receiving simultaneously on all four pairs. Table 1 shows the maximum distances for switched links using the various physical media specified for GigE.

Table 1.—Maximum distances supported by Gigabit Ethernet switched links.

Standard	**Notation**	**Medium**	**Distance (meters)**
IEEE 802.3z (1000BASE-X)	1000BASE-LX	SMF	5,000
		50 μm MMF	550
		62.5 μm MMF	550
	1000BASE-SX	50 μm MMF	500/550
		62.5 μm MMF	220/275
	1000BASE-CX	Shielded Balanced Copper	25
IEEE 802.3ab	1000BASE-T	Category 5 UTP	100

1000BASE-T uses a new 5-level Pulse Amplitude Modulation (PAM) coding scheme and requires the GMII to decouple 8B/10B encoding. In the 5-level PAM coding scheme, four levels are used to represent two bits of information and the fifth is used for Forward Error Correction (Gigabit Ethernet Alliance, March 1999; Gigabit Ethernet Alliance, 1997; IEEE 802.3ab, 1999). The combination of parallel transmission over 4 pairs together with 5-level coding allows 1000BASE-T to send one byte during each signal pulse and, by using a signaling rate of 125 Mbaud, to achieve a date rate of 1 Gbps.

Factors such as signal attenuation, echo, return loss, and crosstalk presented several design challenges to transmitting data at 1 Gbps over four pairs of Category 5 UTP. Each pair is affected by crosstalk from the adjacent three pairs and, while return loss and Far-End Crosstalk (FEXT) have negligible impact when a Category 5 link is used to carry 10BASE-T signals, they can significantly affect the operation of 100BASE-TX and 1000BASE-T LANs. Consequently, in addition to the performance criteria for Category 5 cabling specified in ANSI/TIA/EIA-568A, IEEE 802.3ab has specified new return loss and FEXT tests for 1000BASE-T links. If a link fails to pass any of these tests, the failure is probably due to problems in the connectors or patch cable and corrective actions should be taken.

Upon link initialization, GigE uses an Auto-Negotiation function, which is managed by the PCS sublayer, to negotiate the optimal common mode of operation (Seifert, 1998). Two forms of auto-negotiation are provided for GigE, one for 1000BASE-X networks and the other for 1000BASE-T. As with 10/100 Mbps Auto-Negotiation, devices on a 1000BASE-X link exchange configuration information to determine their modes of operation (i.e., HDX or FDX) and the methods of flow control, if any, that they support. Based on this information, they automatically configure themselves for HDX or FDX operation and, if they support flow control, for symmetric or asymmetric operation, as well as for the direction if asymmetric. However, unlike 10/100 Mbps Auto-Negotiation, 1000BASE-X auto-negotiation does not determine data rate and is restricted to gigabit operation. On the other hand, 1000BASE-T uses the same UTP Auto-Negotiation system employed by 100BASE-TX, extended to include negotiation of the gigabit data rate itself, and is backward compatible with 10BASE-T and 100BASE-T networks.

There are a variety of network applications associated with the various physical layers that GigE supports. 1000BASE-LX is appropriate for inter-building campus backbones while 1000BASE-SX is targeted at shorter intra-building backbone applications and for direct connections to high performance workstations and servers. 1000BASE-CX may be used for short-haul interconnections such as jumper cables in an equipment rack, a server cluster, a wiring closet, or a computer room. 1000BASE-T is intended to take advantage of the extensive installed base of structured Category 5 UTP cabling and may be used to upgrade 10BASE-T and 100BASE-TX links. Like 1000BASE-SX, it is suitable for horizontal cabling on the floor of a building, for high-end desktop computing, and for use within an equipment room or server farm.

3. IEEE 802.3x Flow Control

On shared Ethernet LANs, which are inherently HDX, CSMA/CD acts as a simple flow-control mechanism by preventing more than one station from transmitting at a time. GigE LANs transmit data at 1 Gbps, which is 100 times faster than 10 Mbps Ethernet, and most are switched, with CSMA/CD disabled. Furthermore, GigE switches usually work in FDX mode, in which each connection has the capacity to transmit and receive a total of 2 Gbps, or approx. 1.5 million 64 byte frames per second in each direction. A switch has finite input buffering capacity per port and, if it cannot forward traffic at a rate that is faster than the arrival rate, may become congested during periods of heavy traffic (Kadambi et al., 1998; Seifert, 1998). For example, a server could overwhelm a switch's input buffers by transmitting bursty traffic to a single port on the switch. This is particularly serious if the switch is connected to a lower speed network or if the destination is on a shared segment. If the switch cannot keep up with what it receives from the server, the input port's buffers may overflow. This can cause frame loss and dramatically reduce throughput, especially for applications using protocols such as TCP/IP that eventually retransmit lost packets.

If there is a chance that a switch can become overloaded, flow control becomes critical in order to improve throughput. It can reduce congestion at the link level and prevent buffer overflows and frame loss. It also lowers switch cost by reducing buffer capacity requirements (Seifert, 1998). Gigabit Ethernet switches use the same data link layer Xon/Xoff "stop-start" flow control protocol that was defined in the IEEE 802.3x standard as an option for FDX operation on Ethernet and Fast Ethernet networks (IEEE 802.3, 1998). When an input buffer in a receiving switch is close to capacity (i.e., during periods of congestion), the MAC controller associated with that input buffer sends a "Pause frame" to the source of the congestion. The Pause frame contains a timer value, that is set to the estimated time it will take the congestion to abate, and tells the sending station to stop transmitting for the specified time period. The congested buffer may increase the pause period by issuing another Pause frame before the first period expires and, during the pause period, it can forward queued frames to free up capacity. Once the input buffer's congestion has alleviated, the MAC controller either transmits another Pause frame with a timer value of zero or relies on the expiration of the previous timer value. The sending station is then permitted to resume frame transmission.

Some hubs need asymmetric flow control, which works in only one direction on a link (Kadambi et al., 1998). For example, if an end-station is connected to a hub, the hub can apply flow control to the station, but not vice versa. End-stations are the computers (PCs and workstations) and servers that run network applications and the true sources and sinks of most network traffic. If an end-station could tell a hub to stop transmitting, the hub would stop sending to all attached nodes and bring down that segment of the network. In general, it is desirable to have symmetric flow control for switch-to-switch connections and for a switch connected to a Buffered Distributor (BD, to be discussed in section 4) and asymmetric flow control for a connection from a switch or BD to an end-station. IEEE 802.3x implements both symmetric and asymmetric flow control.

While IEEE 802.3x flow control is simple to implement, it may be too slow to be effective. It is intended to be a low-level scheme for transient congestion and works best on connections between switches and end-stations on small LANs. It does not provide end-to-end flow control (Pause frames are not forwarded by internetworking devices) and is not sophisticated enough for switch-to-switch links on larger networks where the effects of congestion could spread to uncongested segments. For example, switches could propagate jamming signals onto uncongested network segments, preventing users from sending data and creating a congestion effect on segments that have ample capacity. Also, IEEE 802.3x does not distinguish between application streams (it stops all data traffic) or differentiate priorities (which is an issue for delay sensitive traffic) and its implications for higher layer protocol performance such as TCP/IP are not clear. A more sophisticated method of flow control, such as a credit-based or rate-based scheme that can respond with different rates instead of just on and off, may be needed to solve long-term congestion problems. Despite its limitations, IEEE 802.3x is the only standard flow control mechanism defined for FDX Ethernet networks.

4. Switches, Repeaters, and Buffered Distributors

Gigabit Ethernet hubs include switches, repeaters, and new devices called buffered distributors or FDX repeaters. A GigE switch can accommodate dedicated 1 Gbps connections and allow multiple connected stations to transmit simultaneously. Each port includes a GigE physical layer, GigE MAC layer, and input and output buffers. While switch ports may operate in HDX mode, most GigE switches run in FDX mode in which they can provide a combined send-and-receive capability of 2 Gbps per port. A switch may also provide autosensing (for 10/100/1000 Mbps operation) on a port-by-port basis to allow a gradual installation of 100 and 1,000 Mbps devices without an entire network upgrade.

A Gigabit Ethernet repeater is a HDX physical layer device that interconnects Ethernet segments and allows them to share 1 Gbps. All members of the shared network contend for transmission onto a single collision domain and at most one successful transmission is possible at a time. The repeater repeats, or forwards, all incoming frames to all connected ports, except the port on which the frames entered. If it simultaneously detects multiple incoming bit streams, it propagates a jam sequence onto all ports to notify them that a collision has taken place. A repeater does not store frames or have a MAC layer. All ports must operate at the same speed, but they can be connected to any of the standard GigE physical media as long as they use the same encoding method. Due to the bit budget requirements imposed by the CSMA/CD protocol at 1 Gbps, only one repeater is allowed per collision domain. Two or more collision domains may be interconnected with a bridge, switch, or router.

Instead of a HDX repeater, most hub vendors offer a new class of device called a “buffered distributor” or FDX repeater, which is not included in the IEEE 802.3z standard (Kadambi et al., 1998). A BD is a FDX, multiport, hub-like device with multiple GigE ports. It may be used to interconnect two or more IEEE 802.3 links operating at 1 Gbps and to aggregate GigE stations. A buffered distributor combines features found on IEEE 802.3 repeaters and switches. Like a repeater, it is a non-address-filtering device that forwards each incoming frame to all connected links except the originating link. In this manner, similarly to an IEEE 802.3 collision domain, it provides 1 Gbps shared bandwidth to its attached ports. Like a switch, a buffered distributor provides a dedicated point-to-point GigE link to every attached station and is a store and forward device that can simultaneously receive on multiple ports. It eliminates collisions and does not require carrier extension. Also, whereas a conventional repeater is strictly a physical layer device, each port on a BD, as on a GigE switch, includes a GigE physical layer, GigE MAC layer, and input and output buffers.

When an incoming frame enters an input port on a BD, it waits in that port's input queue until it is selected for transmission. Once this occurs, the selected port forwards the frame. If frames arrive at multiple input ports simultaneously, a forwarding protocol (such as round robin) is used to sequentially repeat frames from input ports to output ports. However, the aggregate input rate will equal the number of receiving ports times the 1 Gbps line rate, which exceeds the BD's 1 Gbps total output capacity. Consequently, just as with a switch, the input buffers may become congested and, to prevent frame loss, the BD supports IEEE 802.3x flow control.

Of the three types of GigE hubs discussed above, a GigE switch supports the highest throughputs and the longest distances. It is capable of forwarding 1Gbps per port and supporting distance limits of up to 550 meters over multimode fiber and 5 km over single mode fiber. Also, a switch may incorporate wire-speed forwarding, Virtual LAN (VLAN) tagging, traffic classification, and sophisticated network management capabilities. A large GigE should use FDX switched connections, especially if network expandability is a concern. A GigE repeater, on the other hand, provides the least expensive and least complex method for interconnecting GigE NICs. However, its use of CSMA/CD reduces throughput to less than 1 Gbps and limits network diameter to 200-meters. In addition, repeaters typically have fewer ports than switches and they do not support flow control or VLANs. A BD is cheaper than a GigE switch and has the same distance limits, but it is more expensive than a repeater. Its performance is significantly better than that of a shared, HDX repeater, but considerably inferior to that a switch. Through its use of input buffers and round robin scheduling, a BD can achieve nearly 100% throughput and forward close to 1 Gbps of traffic.

5. Quality of Service Issues and Related Protocols

Network applications such as voice, video, multimedia, and real-time process control have strict QoS requirements, including guaranteed bandwidth and bounds on transmission latency and reception jitter (i.e., delay variance) (Kadambi et al., 1998). For a network to deliver a specific QoS to a particular traffic flow, network switches and routers must set aside resources for that flow. However, since Ethernet was originally intended only to carry data, it included no provisions to provide the QoS guarantees needed by delay-sensitive applications. CSMA/CD simply applies the same access rules equally to all nodes on a shared network, and IEEE 802.3x flow control stops all traffic on a congested, switched network. While it could be argued that GigE offers inherent QoS simply because of its high bandwidth, this would not be valid if congestion is severe in some portion of the transmission path. Therefore, new techniques such as IEEE 802.1 p and Q have been developed to provide Class of Service (CoS) transmission for switched LANs and, at higher layers, some routers implement IP's Type of Service (ToS) field, IP's Differentiated Services (DiffServ), or the IETF's Resource Reservation Protocol (RSVP) (Cunningham et al., 1999). CoS provides a simple traffic prioritization capability which allows frames to be forwarded by network nodes according to their priority levels instead of the order in which they were received. However, unlike QoS, it does not provide guaranteed delivery.

IEEE 802.1p is an extension of the IEEE 802.1D standard for bridging and LAN interconnection (ISO/IEC 15802-3, 1998). It defines how traffic prioritization should be implemented within a MAC-layer bridge (i.e., switch) for Ethernet and other LAN topologies that do not already support priorities. It is a signaling scheme that allows end-stations to request priorities (i.e., classes of service) desired for frames and to communicate these requests to switches along the path. IEEE 802.1p uses a 3-bit "user priority" tag which can be inserted into a MAC frame and whose possible values range from 0 (no priority) to 7 (highest priority). End-stations can set these priority bits to indicate the requested transmission priority level and IEEE 802.1p-compliant switches give higher priority frames precedence for transmission over lower priority or non-tagged frames. Being a Layer 2 mechanism, IEEE 802.1p works on both IP and non-IP networks. However, since the Layer 2 header is only read at the switch level, the boundary routers, where bottlenecks occur, cannot take advantage of IEEE 802.1p unless it is mapped to a Layer 3 prioritization scheme (e.g., IP's ToS maps directly to and from IEEE 802.1p CoS). Thus, while prioritization is accomplished within the switched network, it may be lost at the LAN/WAN boundary. Also, implementing IEEE 802.1p in networks with non-IEEE 802.1p switches that misinterpret the unexpected bits could lead to instability. The non-IEEE 802.1p-compliant devices may interpret a frame as oversized and discard it, or they may pass the frame without the benefit of prioritization.

IEEE 802.1Q is another extension to the IEEE 802.1D bridging standard and defines a general-purpose VLAN implementation (IEEE 802.1Q, 1998). It is closely related to IEEE 802.3ac, which defines the method of VLAN tagging to be used by IEEE 802.3 LANs (IEEE 802.3ac, 1998). VLANs provide a mechanism for classifying traffic and enable a network manager to logically group end-stations into different broadcast domains (Cunningham et al., 1999; Kadambi et al., 1998; Seifert, 1998). Each VLAN is a logical broadcast domain and the end-stations in the same VLAN are able to communicate as if they are physically connected to the same LAN segment, even though they may not be. VLANs are useful for breaking up large Layer 2 LANs into smaller segments and for preventing broadcast storms from overwhelming large switched networks. They also aid routing and management in an extended network. For example, by treating a VLAN identifier as a group address, intermediate network nodes only need to know which ports are associated with each VLAN and thus have fewer address table entries to maintain. Through the use of network management software, VLANs can also simplify moves, adds, and changes in network configuration. Bridging software is used to define which stations are to be included in each VLAN and routers are required for communication between VLANs.

IEEE 802.1Q supports port-based VLAN membership, which allows ports on different switches to be grouped onto the same VLAN and end-station addresses to be associated with VLANs rather than station port numbers (IEEE 802.1Q, 1998). It uses explicit tagging in which the sender's local switch inserts 32 additional bits of data into a MAC frame's header. This 32-bit IEEE 802.1Q header, called a VLAN tag, consists of a Tag Protocol Identifier (TPI) field and a Tag Control Information (TCI) field, as shown in Figure 4.

16	3	1	12
TPI	P	CFI	VI

<------------------ TCI ------------------>

Figure 4.—VLAN tag format.

The TPI field indicates that the frame contains IEEE 802.1Q data and contains the hexadecimal value 81-00, which used to be the Ethertype value for “Welfleet.” The TCI field is divided into three subfields. The P bits contain the IEEE 802.1p user priority value, the Canonical Format Indicator (CFI) bit is set to 0 and not used in IEEE 802.3 networks, and the VLAN Identifier (VI) field indicates the VLAN to which the frame belongs. On IEEE 802.3 networks, the VLAN tag is inserted into a MAC frame’s header between the Source Address field and Type/Length field (i.e., the TPI field is in the location occupied by the Type/Length field in a non-tagged frame). The VLAN tag requires the CRC to be recomputed at its insertion and removal and increases the maximum frame length to 1,522 bytes. The remaining fields in a VLAN tagged frame are the same as in an untagged MAC frame, except that the minimum length of the Data+Pad field is 4 bytes shorter. Once frames are tagged, they can be sent through the network, and through non-IEEE 802.1Q switches that can accommodate the larger frame size, as if they were normal traffic. However, legacy Ethernet devices that participate in VLAN services require new Ethernet cards and software drivers to support the tagged frame format.

Higher layer protocols, of which DiffServ and RSVP are among the most promising, also have a role to play in providing QoS to GigE networks. DiffServ is an IETF QoS standard that operates at Layer 3 (IETF RFC 2475, 1998). It utilizes the Type of Service (ToS) byte in IPv4, or the Traffic Class (TC) byte in IPv6, to mark a packet to receive a particular forwarding treatment or Per-Hop Behavior (PHB) at each network node. Although ToS in the IPv4 header has been available for some time, it has generally been ignored in practice. The DiffServ architecture aims to build a standardized framework in which inter-domain interoperability can be achieved to provide end-to-end QoS. The IPv4 ToS byte (and the IPv6 TC byte) has been renamed the DS byte. By marking the DS field in each packet with a specific value, users can specify the PHB to be allotted to the packet. A PHB, the key building block of DiffServ, defines how traffic belonging to a particular behavior aggregate (i.e., an accumulation of similarly marked packets) is treated at an individual network node. The aggregation of a multitude of QoS-enabled flows into a small number of aggregates, combined with the implementation of complex classification and conditioning functions at network boundary devices, makes DiffServ ideal for deployment in a very large network, such as the Internet, that requires scalability.

RSVP is a layer 4 protocol, also from the IETF, that allows hosts to request specific QoS for application data streams and works with IP to set up communication paths (Cunningham et al., 1999; Kadambi et al., 1998). It is a simple hop-by-hop signaling system in which control packets carry a resource reservation request from a source host through the network. At each router (or other Layer 3 device) on the path to the destination host, RSVP uses admission control to determine if the router has sufficient resources available to satisfy the request and it uses policy control to determine if the user has the administrative permission to make the reservation. If either test fails, RSVP notifies the source that the requested level of service cannot be supported at the present time. Otherwise, RSVP reserves bandwidth (BW) from the router. RSVP requires each network component in the communication path to support RSVP and maintain bandwidth allocation information for each data stream. This can place a heavy load on network resources and reduce capacity for other traffic, especially on large networks. RSVP is being extended to include mechanisms for mapping data streams to IEEE 802.1p and DiffServ service classes. Also, support is growing for a model in which RSVP in the LAN would be integrated with DiffServ in the WAN to achieve end-to-end QoS. In this model, RSVP would negotiate BW reservation at the edge of a network and a border router would map the RSVP parameters to an appropriate DiffServ class for use in the WAN. The benefits include granular QoS at the network edge where specific applications require guaranteed BW, and simpler QoS in the core of the network where scalability and low overhead are needed.

6. Network Management and Performance

IEEE 802.3u Clause 30, which defines the entire management capabilities for 100BASE-T internetworking devices, was taken directly into IEEE 802.3z and enhanced to support the management of GigE as well as integrated 10/100/1000 Mbps Ethernet networks (IEEE 802.3, 1998). IEEE 802.3z has added a number of objects to various categories to support new capabilities created by gigabit operation. Examples include additions to MAC objects (for carrier extension and frame bursting), repeater objects, and MAU types.

Although Clause 30 provides an extensive set of definitions of managed objects for IEEE 802.3 networks, the de facto standard for network management is Simple Network Management Protocol (SNMP) from the IETF (Gigabit Ethernet Alliance, May 1999; Kadambi et al., 1998; Seifert, 1998). A network administrator can use SNMP to view the status of network elements from a central station and use Remote Monitoring (RMON) agents to capture information and send it back to the central console to be analyzed. Most of the work done for Clause 30 has formed the basis for the definition of SNMP Management Information Bases (MIBs). In SNMP terminology, a MIB is a collection of managed objects relating to a specific entity. It specifies various counters, status events, alarms, notifications, and so on, relating to a specific managed device such as a repeater. MIBs, which are different for different devices, are used by SNMP to record statistics such as collision counts, frames transmitted or received, and error rates. Although most GigE switch vendors typically provide some network management capability, and some also provide proprietary extensions to standard MIBs to manage implementation specific features, management of IEEE 802.3 devices is optional and not required for conformance to the standard.

Whether the added capacity of a GigE link provides significant benefit over Fast Ethernet depends on the applications and connected devices. Upgrading connections to GigE in high-end enterprise servers, that typically process data at hundreds of Mbps, should improve utilization. However, replacing a Fast Ethernet connection with GigE could actually reduce throughput for a server that cannot process data at gigabit speeds, or if the server has small memory caches or slow main memories. For example, if a server whose maximum throughput is 100 Mbps is linked in a gigabit connection, the server could be overwhelmed and the resulting lost data and retransmissions could degrade network performance. If IEEE 802.3x flow control is used to prevent a station from sending traffic, delay will increase. However, this is usually preferable to allowing an application to send its traffic and then forcing that traffic to wait in intermediate switch buffers, or to incur the penalty of frame loss due to buffer overflow. In general, wire speed operation is a more critical issue for a campus switch than a workgroup switch and it is important for internetworking devices to forward traffic in increments that are small enough for downstream devices to handle. In addition, only end-stations with high performance processing capabilities and buffer memory will benefit from GigE connections and it only makes sense to increase link capacity if the link itself is the performance bottleneck.

As applications such as intranets and Internet web browsing have increased traffic between subnets, Layer 3 internetworking devices have become potential performance bottlenecks in GigE networks (Kadambi et al., 1998). They perform complex tasks such as packet conversion, segmentation and reassembly, and encapsulation and decapsulation. For example, a router typically examines the destination address field in each incoming packet, uses the address as an index into a routing table to determine the next hop, and then modifies and copies the packet to the output interface. These tasks, which have traditionally been implemented in software, increase processing overhead. Many GigE switches now include new built-in Layer 3 switching techniques that accomplish these tasks more quickly and efficiently in hardware. They generally support IP and some also implement other Layer 3 protocols such as Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Next Hop Resolution Protocol (NHRP) (Cunningham et al., 1999). IP switching uses the concept of a flow. This is a sequence of packets that are forwarded, from a source to a destination, through the same ports and with the same QoS. A flow classification, which determines how further packets belonging to the flow are forwarded, is selected based on the first few packets in the flow. The classification is cached and therefore does not require a full time-consuming lookup for every packet.

Internal architecture is a key consideration for the performance of Gigabit Ethernet hubs. The internal BW of a repeater or BD needs to be no greater than that of a single port (i.e., 1 Gbps) while the internal BW of a switch must be greater than the individual port capacities (i.e., multiple Gbps) (Seifert, 1998). If a switch uses a blocking architecture, it will not be able to support traffic patterns in which all links carry sustained traffic at full speed without either discarding frames or invoking flow control. Depending on load patterns, the higher cost of nonblocking may be justified in order to enable wire-speed forwarding and prevent packets from being discarded. A nonblocking GigE switch requires a backplane capacity that equals or exceeds the total capacities of its input ports. For example, a nonblocking FDX switch configured with ten 100 Mbps ports and one 1 Gbps port would need an internal bus with a capacity of at least 4 Gbps. However, there could still be congestion due to output port blocking if there is more traffic destined for a given output port than the port can handle.

For shared networks, GigE has a longer slot time than 10 and 100 Mbps Ethernet and this reduces efficiency. Due to the higher ratio of round trip propagation delay to frame transmission time, especially for short frames, a higher proportion of time is spent in collision resolution than in frame transmission. In addition, carrier extension further degrades throughput for frames shorter than 512 bytes, which may have up to 448 bytes of padding. For example, in the worst case, the channel efficiency for a stream of 64 byte frames with 64 bit preamble/start-of-frame delimiter and 96 bit IFG is 512/(4096+64+96) or 12%, versus 512/(512+64+96) or 76% for 10 and 100 Mbps networks (Seifert, 1998). In general, the distribution of frame sizes being carried has a significant impact on GigE performance and, when sending a large number of small frames, the throughput is only marginally better than Fast Ethernet. While frame bursting can improve this situation, most applications cannot take advantage of the technique.

7. Gigabit Ethernet versus ATM

ATM, which like GigE operates at the data link and physical layers of the OSI model, is the other dominant technology competing for use in building and campus backbone networks (Bakes et al., 1996). It is a connection-oriented, fast packet switching technology that uses 53 byte cells to transport information and statistical multiplexing to provide bandwidth on demand. The 53-byte cell size was selected as a compromise between the goals of packetization delay and payload efficiency. Short fixed length cells reduce delay and jitter and are therefore appropriate for transporting delay-sensitive voice and video traffic, whereas long cells have lower overhead and are more efficient for data applications. Each cell consists of a 5-byte header and 48-byte payload (Figure 5).

Header	Payload
<----- 5 ----->	<------------------------------- 48 ------------------------------->

Figure 5.—ATM cell format.

ATM is able to offer true QoS based on performance parameters that are negotiated across a User-Network Interface (UNI) between an attached station and ATM switch prior to sending user information. The station uses UNI signaling to request a certain QoS level for each application and, if the network can guarantee the requested QoS, a virtual path is established to support the application. Otherwise the connection is refused. UNI specifications define an explicit set of performance parameters such as maximum Cell Transfer Delay (i.e., latency), peak-to-peak Cell Delay Variation (i.e., jitter), and cell loss ratios. Based on these parameters, ATM QoS is classified into service classes (Bakes et al., 1996; Kadambi et al., 1998). AAL1 is used to support real-time constant bit rate traffic such as voice and video, AAL2 is used to support real-time VBR traffic such as MPEG video, and AAL 3/4 is used to support non-real-time data. AAL 3/4 was originally intended to carry LAN traffic but, for this purpose, has since been replaced by AAL5. These traffic classes map to four traffic types, which are referred to as Constant, Variable, Available, and Unspecified Bit Rates (i.e., CBR, VBR, ABR, and UBR). Due to its use of small fixed length cells and QoS parameters, ATM can carry voice, video, data, imaging, and graphics, separately or simultaneously, on the same link. ATM networks also implement sophisticated credit and rate-based congestion control and support Private Network-to-Network Interface (PNNI), which is a QoS-aware routing protocol.

For existing connectionless protocols such as IP and Ethernet to work over ATM, they must be adapted to operate directly over an AAL via Classical IP over ATM (CIP), ATM LAN Emulation (LANE), or Multiprotocol Over ATM (MPOA). For example, an Ethernet device can use LANE to pass Ethernet MAC frames over an ATM network. LANE is a Layer 2 protocol and CIP and MPOA operate at Layer 3.

CIP, as specified in RFC 1477, allows IP traffic to be routed over an ATM network and is transparent to the TCP/IP stack (Cunningham et al., 1999). It enables an ATM attached device to transmit IP packets and communicate with an IP device. ATM Address Resolution Protocol (ATMARP) and inverse ATMARP are used to map IP addresses to and from ATM addresses, respectively. CIP is based on the concept of a Logical IP Subnetwork (LIS) which contains hosts and routers having the same IP subnet mask and same subnet address. Hosts in the same LIS communicate directly using virtual channels and hosts from different LISs communicate through a router. However, CIP has no support for multicast traffic and IP is the only protocol that it can run. Also, ATM's intrinsic QoS properties may be lost passing through routers.

The ATM Forum created the LANE specification to enable legacy LAN (e.g., Ethernet) devices to interoperate across an ATM network and with devices that connect directly to ATM switches (Kadambi et al., 1998). For example, LANE allows devices on different Ethernet segments to communicate with one another across an ATM network in a way that makes the ATM network transparent to the Ethernet segments, upper layer protocols, and end user applications. Each legacy LAN requires a device called a LEC, attached between it and the ATM network, to act as a bridge and convert MAC frames to and from ATM cells. The ATM Forum specifies AAL5 as the AAL for use with LANE. LANE provides permanent and switched virtual circuit connections, uses its Broadcast and Unknown Server (BUS) capability to support multicast traffic, and allows the creation of VLANs. However, since LANE emulates a traditional LAN interface, it cannot take advantage of ATM's native QoS features.

MPOA, also from the ATM Forum, enables applications with different network layer protocols (e.g., IP-based applications) and underlying networks (e.g. Ethernet) to be routed and bridged across an ATM network (Cunningham et al., 1999; Kadambi et al., 1998). While LANE makes ATM transparent to Layer 2 networks, MPOA also makes ATM transparent to Layer 3 networks. MPOA provides end-to-end network layer connectivity and virtual routing for hosts that are directly attached to the ATM network or indirectly attached from a legacy LAN IP subnet. It transports each traffic flow, from source to final destination, via a single-hop switched virtual circuit connection (called a shortcut ATM path) and makes the entire ATM network appear as one logical router hop. MPOA uses NHRP to determine the shortcut paths and LANE for bridging and configuration purposes. It supports all of ATM's QoS features, but legacy LAN devices may not be able to take advantage of these.

As explained in the remainder of this section, GigE and ATM each has its own strengths and limitations.

7.1 Quality of Service

As discussed in section 5, network applications such as voice, video, and multimedia have strict requirements for QoS which, in addition to traffic prioritization, implies a guarantee of bandwidth and bounds on latency, jitter, and error rate. While it could be argued that Gigabit Ethernet offers inherent QoS simply because of its high bandwidth, classic Ethernet is a data-only transport that does not provide the QoS guarantees needed for delay-sensitive traffic. It is a connectionless technology that transmits variable-length frames. It cannot differentiate between applications or guarantee that real-time traffic gets the preferential treatment it requires and it is possible for a small time-sensitive frame to get delayed behind a large data frame. New techniques such as IEEE 802.1p/Q and RSVP allow CoS capabilities to be implemented on Ethernet LANs by assigning priorities to specific VLANs, end-stations, or application sessions. However, while CoS techniques can be used to prioritize frames, they cannot reserve bandwidth for an entire application stream and are unable to provide guaranteed QoS. Also, being new, they may have interoperability problems with existing infrastructures.

ATM, on the other hand, was designed to deliver true QoS capabilities for high-quality voice and video and supports CBR, VBR, ABR, and UBR traffic types. Unlike Ethernet, it is a connection-oriented scheme that transmits short, fixed-length cells, allows bandwidth to be reserved for an entire stream, and guarantees a constant level of service for the duration of a session. However, while ATM can implement LANE to support VLANs, only native ATM is able to offer guaranteed QoS features. Also, in order to establish an ATM connection, an application must know its communications requirements in advance, which may be reasonable for voice/video services but not for computer data.

7.2 Data Rate and Throughput

Ethernet is scalable from 10 to 100 to 1,000 Mbps (and 10 Gbps is under investigation), which allows an incremental migration to higher-speed networking and is important for LAN backbones that have become congested. A consistent Ethernet environment avoids the performance penalties for the frame and media conversions that are normally required when translating between different LAN types. To carry traffic from higher level protocols, both GigE and ATM must encapsulate the higher level packets, which typically requires less overhead with GigE than with ATM (Kadambi et al., 1998). For example, in the case of a 1500 byte IP packet, GigE adds 26 bytes of overhead and transmits a total of 1526 bytes. ATM AAL5, on the other hand, adds an 8 byte trailer plus a 28 byte pad to ensure the AAL5 Protocol Data Unit (PDU) is a multiple of 48 bytes. The resulting 1,536 byte AAL5 PDU is then divided into 48 byte segments and transmitted in 32 ATM cells, each with a 5-byte header and 48-byte payload, for a total of 1696 bytes. Thus, the added overhead required to transmit the IP packet is only 2% with GigE versus 12% with ATM. However, as explained in section 6, the requirement for carrier extension on shared GigE networks causes inefficient use of bandwidth and reduces throughput, especially when sending small frames. In addition, if an Ethernet LAN is to be connected to an ATM WAN, a switch or router is required to translate Ethernet frames to or from ATM cells and these conversions reduce effective throughput. Furthermore, on FDX Ethernet networks, IEEE 802.3x flow control is the only standard flow control mechanism which, while adequate as a low-level scheme for transient congestion on small LANs, may not be able to solve long-term congestion problems in large LANs.

ATM is also scalable and generally uses a SONET physical layer. ATM links are capable of operating at a wide range of data rates, including sub-T1, 1.544 Mbps [T1], 25 Mbps, 155 Mbps [OC-3], 622 Mbps [OC-12], 2.4 Gbps [OC-48], and 10 Gbps [OC-192] (Bakes et al., 1995). ATM's use of small fixed size cells enables fast and efficient hardware implementations of ATM switches. It also allows memory to be allocated in exact increments, which reduces wasted storage and allows efficient address lookup. Unlike GigE, for which the distance limits and MAC layer implementation are different for different data rates, ATM is independent of data rate and physical layer technology. In addition, implementing ATM in both the LAN and WAN environments avoids having to translate frames to or from cells, which improves throughput and latency. Furthermore, ATM networks are able to implement sophisticated credit and rate-based congestion control schemes. However, short cells require more cells for a given amount of information, which increases overhead for headers and processing requirements at switches. Also, an ATM switch generally allows for an occasional cell to be discarded under congestion conditions, which can cause severe degradation in performance for data communications applications. Following the loss of a single cell, a higher layer will implement an error control mechanism that could involve retransmitting the entire network layer packet, or even the entire window of packets.

7.3 LAN/WAN Scalability

Gigabit Ethernet is essentially a "campus technology." It is primarily intended for use as a backbone and to connect servers, server farms, and powerful workstations in a campus-wide network. With the exception of 1000 BASE-LX over SMF which extends to 5 km, GigE implementations have maximum distance specifications of 550 meters or less which limits their use in Metropolitan Area Networks (MANs) and prevents their use in WANs. Also, QoS capabilities implemented on top of Ethernet are unlikely to scale well in large enterprise environments.

ATM has no physical media distance limits. It can scale from the desktop to host servers to the LAN or campus backbone, and from LAN to WAN, all under a single architecture. It offers seamless LAN/WAN interconnection and, unlike GigE, can be used to provide WAN access and transport services.

7.4 Routing and Addressing

Since GigE uses the same IEEE 802.2 Logical Link Control (LLC) as standard Ethernet, existing network protocols such as IP and Internet Packet Exchange (IPX) operate over GigE without modification. In addition, built-in Layer 3 switching is available with many GigE devices to provide wire-speed routing and simplified packet processing. Most GigE devices support IP and some also support IPX.

ATM devices also provide wire-speed routing and support multiple traffic protocols. Many commonly used data communication protocols, such as Ethernet and IP, are connectionless and rely heavily on broadcast and multicast capabilities for functions such as address discovery, service advertisements, and routing table updates. However, these capabilities are difficult to implement in a connection-oriented ATM network and, as a result, transporting connectionless protocols over ATM generally requires complex higher layer protocols such as classical IP over ATM, ATM LANE, MPOA, or PNNI.

7.5 Interoperability

Upgrading to Gigabit Ethernet is relatively seamless. GigE is compatible with Ethernet and Fast Ethernet and is more likely than ATM to be compatible with installed server, desktop, and network infrastructure equipment. All applications that work on Ethernet will work on GigE. GigE requires no changes to higher layer protocol stacks (such as TCP/IP and IPX), software applications, or operating systems, although it may be appropriate to “tune” the behavior of the upper-layer protocols and applications to take advantage of the increased available BW. On the other hand, the product maturity of GigE is less than that of ATM and, especially for pre-standard products, interoperability among GigE devices from different vendors is an issue. Also, QoS capabilities implemented on top of Ethernet may have interoperability problems across different vendors' equipment.

ATM-based switches have been widely deployed and have proven interoperable in campus backbones, enterprise networks, and private and public WANs. They offer seamless LAN/WAN integration. However, as explained above, running current applications on an ATM network requires protocols such as LANE or MPOA.

7.6 Network Management

GigE provides the same management tasks as 10 and 100 Mbps shared and switched Ethernet networks. As discussed in section 6, IEEE 802.3u Clause 30 from the standard for Fast Ethernet was enhanced to provide network management for 10/100/1000 Mbps integrated Ethernet networks. However, managing switched networks at gigabit data rates is more difficult than at lower data rates and could degrade network performance. Also, GigE has no out-of-band capabilities for enhanced network management and no link fault diagnostics.

ATM's ability to scale from LAN to WAN under a single architecture simplifies network design and management. ATM switches furnish detailed statistics on each connection and each link. It is also possible to monitor standard VLAN-based LANE implementations and MPOA server capabilities, all from a centralized network operation center. ATM technology has F1 to F5 Operations, Administration, and Maintenance (OAM) flows for embedded management, and fault management is available via loopback at different flow levels.

7.7 Cost of Ownership

In general, the total cost of ownership for Gigabit Ethernet can be much lower than for ATM (Gigabit Ethernet Alliance, May 1999). Assuming identical physical media interfaces, GigE is currently cheaper per network adapter and per switch port than a 622 Mbps ATM interface. Furthermore, due to competition and economies of scale, the per-port cost of Ethernet and Fast Ethernet products has decreased significantly in recent years and the cost of GigE interfaces are likely to show similar price declines. Low cost per port is particularly important for desktop connections due to their large numbers. The IEEE's goal is to provide a GigE connection at two to three times the cost of a 100BASE-FX interface.

In addition to the purchase price of the equipment, the total cost of network ownership includes installation, training, maintenance, and troubleshooting costs. GigE networks operate over the same wiring infrastructures as lower data rate Ethernet and Fiber Distributed Data Interface (FDDI) networks. They are relatively easy to install, support, and administer and, due to wide familiarity with Ethernet technology, require minimal new training for support staff. The network operating system, software applications, NIC drivers, and protocol stacks can remain unchanged and only incremental purchases of maintenance and troubleshooting tools are likely to be needed. Furthermore, many GigE devices incorporate Layer 3 switching which is essentially giving away high speed routing. Consequently, GigE networks can usually be deployed more quickly and inexpensively than alternative technologies.

ATM can be deployed in LAN, campus, MAN, and WAN environments under a single architecture. It allows data, voice, and video traffic to be transported over a single integrated network instead of multiple dedicated networks. ATM thus simplifies network design and management, maximizes skill sets and network architecture experience, and can lower total cost of ownership, especially when several geographically dispersed locations are to be interconnected. However, some network managers consider ATM too complex and have concerns about a variety of issues such as the number of switched connections per second a device can handle, LANE compatibility, the state of MPOA standards, and multicast and broadcast traffic.

8. Gigabit Ethernet Applications at the NASA Glenn Research Center

In spite of the relatively short history of Gigabit Ethernet technology and standards, NASA Glenn's research and business communities have actively adopted this promising networking technology into their computing environments. GigE can support a variety of applications, multiple data types, and a large number of users. This has become possible due to a combination of increased bandwidth, LAN switching, protocols such as RSVP that provide bandwidth reservation, standards such as IEEE 802.1p and Q that support packet prioritization and VLANs, and the use of video compression such as MPEG-2.

Many network applications involve high-resolution graphics, real-time video, and other multimedia data types that can benefit from the high bandwidth provided by Gigabit Ethernet (Gigabit Ethernet Alliance, May 1999). For example, engineers and scientists often work interactively in distributed development teams, using design automation tools, interactive whiteboarding, file sharing, and desktop videoconferencing. In such situations, GigE can support multiprocessor applications and expedite the transfer of large Computer-Aided Design/Computer-Aided Manufacturing (CAD/CAM) files or 3-D visualizations of aircraft and other complex objects. GigE can also be used in private Intranets to carry text, graphics, and images, as well as more bandwidth-intensive audio, video, and voice traffic, and in data warehouse applications where large quantities of data are distributed over many platforms, accessed by a large number of users, and regularly updated. In addition, network backups of enterprise information require large amounts of bandwidth for fixed amounts of time. These usually occur during off-hours (e.g., overnight) and involve up to terabytes of data distributed over hundreds of servers and storage systems.

Advances in processor, memory, and disk storage technologies, combined with the availability of high speed networking, have led to the emergence of distributed, workstation clusters as powerful, low-cost alternatives to conventional supercomputing systems for scientific computing applications. The Advanced Computational Concepts Laboratory (ACCL) at NASA Glenn provides an affordable, high performance, multi-platform, computing environment for Glenn's researchers. The platforms are typically characterized by high speed multi-processors, enhanced memory and graphics cards, and advanced networks. ACCL also houses testbeds for exploration of emerging network and computing technologies. There is a LINUX-based, parallel testbed that consists of 32 data nodes and 8 router nodes which are interconnected using Gigabit Ethernet technology configured in a 2-level tree topology (Sang et al., 1999). The router nodes have Pentium II 400 MHz single processors and are interconnected by a Gigabit Ethernet buffered distributor. The data nodes are connected to the router nodes via Fast Ethernet and have Pentium II 400 MHz dual processors. The much improved throughput in inter-processor communication due to these high speed network connections brings an enormous performance benefit to computationally intensive applications. Another testbed, a 24-processor SGI cluster, supports the Information Power Grid (IPG) project in collaboration with the NASA Langley and Ames Research Centers. The goal of the IPG project is for NASA researchers to be able to initiate a process from any of the three Centers and, depending upon where specified resources are available, the process can be scheduled to execute via various job schedules. The SGI cluster currently uses Fast Ethernet for inter-process communication, but upgrading to Gigabit Ethernet would improve performance.

High-performance computing, combined with advanced networking technologies, enables the modeling and simulation of an entire aircraft engine system. Due to limited computing resources, NASA Glenn researchers have traditionally performed aerodynamic and thermal analysis for each engine component separately. The drawback to this approach is that much of the detailed flow physics at the interface between two components can be lost. The Numerical Propulsion System Simulation (NPSS) project is a NASA/Industry joint effort to provide the aeropropulsion industry with the ability to perform detailed computer simulations of complete aircraft engines. The high performance computing system software and Common Object Request Broker Architecture-based (CORBA) object technology enable distributed and heterogeneous computing platforms to be linked and to operate as a seamlessly integrated system. Unix workstation clusters, with a mix of Fast Ethernet and ATM connections, currently support the NPSS project. The high throughput advantage of Gigabit Ethernet is expected to boost their performance.

Gigabit Ethernet technology is also used in NASA Glenn's Telescience Support Center (TSC). TSC is a NASA telescience ground facility that provides the capability to execute ground support operations of in-orbit International Space Station and space shuttle payloads. Through TSC, payload developers and scientists can remotely control and monitor their on-board payloads from any location, usually their home sites, which enhances the quality of scientific and technological data while decreasing operations costs. TSC acts as a hub to provide video distribution and recording services, video and voice conferencing services, and high speed networking services to customers. Two Ethernet switches, with a mix of Fast Ethernet and GigE interfaces, support the data systems.

A plan to use Gigabit Ethernet technology to investigate the performance of frame-based layer two protocols in space communication research is underway. The network uses a hybrid satellite/terrestrial topology, OC-12 ATM links from the NASA Advanced Communications Technology Satellite (ACTS) satellite to switches on Earth, and either GigE or ATM links to researchers' workstations. Interoperability tests with frame-based protocols are to be conducted between different operating systems to investigate how the performance of Gigabit Ethernet-ATM-Gigabit Ethernet circuits compares with that of a Gigabit Ethernet-Packet over SONET-Gigabit Ethernet circuits.

The most promising application of Gigabit Ethernet technology at the NASA Glenn Research Center is in the next generation campus backbone network. The existing campus backbone consists of a number of routers interconnected via a 100 Mbps FDDI network (Bakes et al., 1995). FDDI's main role has been as a highly reliable backbone but, while both shared and switched FDDI products are available, there has been no activity to increase the data rate above 100 Mbps. At the edge of NASA Glenn's FDDI network, users at their desktops connect to the backbone through hubs and send and receive data via 10 Mbps shared Ethernet technology. In these TCP/IP-Ethernet hubs, all application types are treated equally and contend for a fixed amount of bandwidth. Although the current network topology at NASA Glenn has adequately supported network applications in the past, problems associated with bandwidth shortage, application bottleneck, and slow response time have been observed. Many emerging applications, such as real-time multimedia conferencing and online distance learning, usually require much higher throughput, as well as minimum delay and better security. Some users want preferential treatment in terms of guaranteed bandwidth and response time for their applications. In order to meet such demands for more bandwidth and proactively plan for improved network services, Glenn's network infrastructure has to be dramatically updated. As a migration strategy to provide more bandwidth to the desktop and satisfy the needs of bandwidth-intensive applications, Glenn has deployed Gigabit aggregators for some power users. Through the Gigabit aggregators, desktops with Fast Ethernet interfaces can achieve 100 Mbps bandwidth end-to-end.

Note that a Gigabit aggregator functions like a concentrator in order to trunk Fast Ethernet traffic from multiple sources. Eventually, the ever-increasing network traffic for intranet and Internet applications, combined with users' demands for better quality and security, is expected to put Glenn's current backbone network infrastructure at risk of becoming obsolete. Switches based on Gigabit Ethernet technology and standard protocols provide ample bandwidth, wire-speed performance, quality of service, multicast capability, and better security. Glenn's advanced network architecture team is actively looking into deploying Gigabit Ethernet switch technology for its next generation backbone network.

When combined with emerging QoS standards, Gigabit Ethernet's high speed and use of packet switching technology provide a very favorable environment for Voice over IP (VoIP) applications. The consolidation of voice and data on a unified network brings many benefits in terms of savings in capital and management costs, reduced staff requirements, and so on. Nevertheless, there are numerous issues to be resolved before end-to-end VoIP can be fully realized. Like most other organizations, Glenn currently has separate infrastructures for voice and data traffic. As the life cycle of the existing Private Branch Exchange-based (PBX) voice system approaches its end, the convergence of voice and data applications over packet switching has to be considered. A group of Gigabit switch manufacturers is also developing a fiber-based solution to extend the distances supported by Gigabit Ethernet to 50-70 km. The use of Gigabit Ethernet technology for this distance range can provide a viable, cost effective alternative to Glenn's WAN connection to the Plumbrook Station which houses key facilities for space power and propulsion experiments at a location about 45 miles from the main Center.

9. Conclusion

As discussed throughout this paper, GigE and ATM each has its own strengths and limitations. Ethernet has evolved from a HDX, shared network to a FDX, switched network, and its data rate has scaled from 10 to 100 to 1,000 Mbps. Gigabit Ethernet, with its data rate of 1 Gbps, is fully compatible with Ethernet and Fast Ethernet and offers seamless migration to even higher speeds, enabling existing networks to be upgraded without having to change their wiring, higher layer protocols, or applications. GigE is appropriate for high speed, connectionless, data applications, where low cost and high throughput are required, and where WAN integration and strict QoS are not primary concerns. GigE is an ideal technology for IP-based data traffic and, by using higher layer protocols such as IEEE 802.1 p/Q, DiffServ, and RSVP, has the ability to provide CoS capabilities for multimedia traffic. It solves the problem of how to reduce delay on congested networks by increasing BW rather than by increasing complexity. While this does not make sense for a WAN environment where BW is expensive, it is highly appropriate for campus and building networks where the cost of higher BW is usually less than the cost of implementing QoS complexity.

In comparison with GigE, ATM provides similar bandwidth, more functionality, and improved QoS, but at higher cost. ATM is robust, scalable in terms of distance as well as data rate, appropriate for use in both LAN and WAN environments, and able to carry voice, video, and other delay-sensitive applications over a single integrated, connection-oriented network. However, while native ATM provides guaranteed QoS that is ideal for voice and video traffic, higher layer protocols are needed to transport IP-based data traffic over ATM and their use may result in a loss of the ability to deliver true QoS. In essence, ATM is optimized for characteristics that are generally irrelevant for data communications applications which tend to be delay insensitive, jitter insensitive, asymmetric, and loss sensitive.

In conclusion, GigE offers a cheaper and simpler solution than ATM for campus and building backbones, as well as for desktop, workgroup, and server connections, in areas where legacy LAN technologies are no longer able to provide adequate BW. It provides a cost-effective solution for upgrading NASA Glenn's aging FDDI network, as well as some of the attached 10 and 100 Mbps Ethernet networks. Gigabit Ethernet offers the high bandwidth and prioritization capabilities required to support mission-critical research and development activities at the NASA Glenn Research Center and is an appropriate technology for the next generation campus backbone.

10. Abbreviations

Abbreviation	**Explanation**
μm	microns
AAL	ATM Adaptation Layer
ABR	Available Bit Rate
ACCL	Advanced Computational Concepts Laboratory
ANSI	American National Standards Institute
ATM	Asynchronous Transfer Mode
BD	Buffered Distributor
BW	Bandwidth
CBR	Constant Bit Rate
CFI	Canonical Format Indicator
CIP	Classical IP over ATM
CoS	Class of service
CRC	Cyclic Redundancy Check
CSMA/CD	Carrier Sense, Multiple Access/Collision Detection
DiffServ	Differentiated Services
DMD	Differential Mode Delay
DS	Differentiated Services
FDDI	Fiber Distributed Data Interface
FDX	Full Duplex
FEXT	Far-End Crosstalk
Gbps	Gigabits per second
GigE	Gigabit Ethernet
GMII	Gigabit Media Independent Interface
HDX	Half Duplex
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IFG	Inter-Frame Gap
IP	Internet Protocol
IPG	Information Power Grid
IPX	Internet Packet Exchange
km	kilometers
LAN	Local Area Network

LANE	LAN Emulation
LED	Light Emitting Diode
LIS	Logical IP Subnetwork
m	meters
MAC	Media Access Control
MAN	Metropolitan Area Network
Mbps	Megabits per second
MDI	Medium Dependent Interface
MIB	Management Information Base
MMF	Multimode Fiber
MPOA	Multiprotocol Over ATM
NHRP	Next Hop Resolution Protocol
NIC	Network Interface Card
nm	nanometers
NPSS	Numerical Propulsion System Simulation
OC	Optical Carrier
OSI	Open Systems Interconnection
PAM	Pulse Amplitude Modulation
PCS	Physical Coding Sublayer
PDU	Protocol Data Unit
PHB	Per Hop Behavior
PMA	Physical Medium Attachment
PMD	Physical Medium Dependent
PNNI	Private Network-to-Network Interface
QoS	Quality of Service
RFC	Request For Comments
RS	Reconciliation Sublayer
RSVP	Resource Reservation Protocol
SFD	Start-of-Frame Delimiter
SMF	Single-Mode Fiber
SNMP	Simple Network Management Protocol
TC	Traffic Class
TCI	Tag Control Information
TCP	Transport Control Protocol
TIA	Telecommunications Industries Association
ToS	Type of Service
TPI	Tag protocol identifier
TSC	Telescience Support Center
UBR	Unspecified Bit Rate
UNI	User-Network Interface
UTP	Unshielded Twisted-Pair
VBR	Variable Bit Rate
VI	VLAN Identifier
VLAN	Virtual LAN
VoIP	Voice over IP
WAN	Wide Area Network

11. References

Bakes, C.M. and Goldberg, F.N., 1996. ATM networks and their applications at the NASA Lewis Research Center: A case study. *Journal of Engineering and Technology Management*, 13: 1–28.

Bakes, C.M. and Goldberg, F.N., 1995. Applications of Fibre Optic Networks in High Technology Research. *International Journal of Computer Applications in Technology*, 8 (3/4): 172–189.

Cunningham, D.G. and Lane, W.G., 1999. *Gigabit Ethernet Networking.* Macmillan Technical Publishing, Indiana, 564 pp.

Gigabit Ethernet Alliance, May 1999. *Whitepaper: Gigabit Ethernet Overview*. **http://www.gigabit-ethernet.org/technology/whitepapers/gige_0698/papers98_toc.html**

Gigabit Ethernet Alliance, March 1999. *Whitepaper: Gigabit Ethernet Over Copper*. **http://www.gigabit-ethernet.org/technology/whitepapers/gige_0399/copper99_toc.html**

Gigabit Ethernet Alliance, 1997. *Whitepaper: Gigabit Ethernet Over Copper*. **http://www.gigabit-ethernet.org/technology/whitepapers/gige_11.97/papers97_toc.html**

IEEE 802.1Q, 1998. *IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridge Local Area Networks.* IEEE, New York, 214 pp.

IEEE 802.3, 1998. *Information technology--Telecommunications and information exchange between systems--Local and metropolitan area networks--Specific requirements--Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.* IEEE, New York, 1268 pp.

IEEE 802.3ab, 1999. *Supplement to IEEE Std 802.3, 1998 Edition: Physical Layer Parameters and Specifications for 1000 Mb/s Operation over 4-Pair of Category 5 Balanced Copper Cabling, Type 1000BASE-T.* IEEE, New York, 144 pp.

IEEE 802.3ac, 1998. *Supplement to IEEE Std 802.3, 1998 Edition: Frame Extensions for Virtual Bridged Local Area Networks (VLAN) Tagging on 802.3 Networks.* IEEE, New York, 20 pp.

IETF RFC 2475, 1998. *An Architecture for Differentiated Services.* **http://www.ietf.org/html.charters/diffserv-charter.html**

ISO/IEC 15802–3, 1998. [ANSI/IEEE Std 802.1D, 1998 Edition] *Information technology--Telecommunications and information exchange between systems--Local and metropolitan area networks--Common Specifications--Media access control (MAC) bridges.* IEEE, New York, 384 pp.

Kadambi, J., Crayford, I. and Kalkunte, M., 1998. *Gigabit Ethernet: Migrating to High-Bandwidth LANs.* Prentice-Hall, Inc., New Jersey, 365 pp.

Sang, J., Kim, C.M, Kollar, T.J. and Lopez, I., 1999. High-Performance Cluster Computing over Gigabit/Fast Ethernet. *Informatica*, 23: 19-27.

Seifert, R., 1998. *Gigabit Ethernet: Technology and Applications for High speed LANs.* Addison-Wesley Longman, Inc., Massachusetts, 411 pp.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY (*Leave blank*)

2. REPORT DATE
February 2000

3. REPORT TYPE AND DATES COVERED
Technical Memorandum

4. TITLE AND SUBTITLE
An Assessment of Gigabit Ethernet Technology and Its Applications at the NASA Glenn Research Center

5. FUNDING NUMBERS
WU–509–10–24–00

6. AUTHOR(S)
Catherine Murphy Bakes, Chan M. Kim, and Calvin T. Ramos

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
John H. Glenn Research Center at Lewis Field
Cleveland, Ohio 44135–3191

8. PERFORMING ORGANIZATION REPORT NUMBER
E–12108

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546–0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA TM—2000-209803

11. SUPPLEMENTARY NOTES
Catherine Murphy Bakes, Department of Administrative Sciences, College of Business Administration, Kent Sate University, Kent, Ohio 44242; Chan M. Kim and Calvin T. Ramos, NASA Glenn Research Center. Responsible person, Chan M. Kim, organization code 7190, (216) 433–8715.

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified - Unlimited
Subject Category: 62 Distribution: Nonstandard

This publication is available from the NASA Center for AeroSpace Information, (301) 621–0390.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 words)***
This paper describes Gigabit Ethernet and its role in supporting R&D programs at NASA Glenn. These programs require an advanced high-speed network capable of transporting multimedia traffic, including real-time visualization, high-resolution graphics, and scientific data. GigE is a 1 Gbps extension to 10 and 100 Mbps Ethernet. The IEEE 802.3z and 802.3ab standards define the MAC layer and 1000BASE-X and 1000BASE-T physical layer specifications for GigE. GigE switches and buffered distributors support IEEE 802.3x flow control. The paper also compares GigE with ATM in terms of quality of service, data rate, throughput, scalability, interoperability, network management, and cost of ownership.

14. SUBJECT TERMS
Gigabit ethernet; Campus backbone; NASA Glenn; Quality of service; ATM

15. NUMBER OF PAGES
33

16. PRICE CODE
A03

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	